CLAIMS

What is claimed is:

1. A method of routing a flow of frames through a switch comprising:

receiving at least one frame from said flow of frames;

applying a process to select an exit port of said switch from a set of possible exit ports through which at least one frame from said flow of frames will exit so as to potentially reduce frame traffic congestion along potential routes that include said set of possible exit ports, said set of possible exit ports including at least some of the exit ports of at least two trunk groups;

transmitting said at least one frame.

2. The method of claim 1, wherein said set of possible exit ports includes at least all of the exit ports of at least two trunk groups.

3. The method of claim 1, wherein at least one of said trunk groups comprises four exit ports.

4. The method of claim 1, wherein at least one of said trunk groups comprises eight exit ports.

5. The method of claim 1, wherein said process comprises a pseudo-random process.

6. The method of claim 5, wherein applying said pseudo-random process comprises applying a hash function.

7. The method of claim 6, wherein said hash function is applied to a set of parameters associated with the frames exiting said switch in order to select an exit port from said set of possible exit ports.

8. The method of claim 1, wherein a weight is respectively assigned to at least some respective ones of said exit ports;

wherein applying a process to select an exit port of said switch from a set of possible exit ports through which a frame from said flow of frames will exit so as to potentially reduce frame traffic congestion along potential routes that include said set of possible exit ports comprises employing the weights to select an exit port that is as good as or better than alternative exit ports in terms of achieving an objective function reflected by said weights.

9. The method of claim 8, wherein as good as or better in terms of achieving an objective function comprises obtaining a higher value objective function.

10. The method of claim 8, wherein as good as or better in terms of achieving an objective function comprises obtaining a lower value objective function.

11. The method of claim 8, wherein said weights at least in part reflect consumed bandwidth for particular routes.

12. The method of claim 8, wherein at least some exit ports have multiple weights reflecting routes from the particular exit port to multiple respective destination ports;

wherein applying a process to select an exit port from a set of possible exit ports through which a frame from said flow of frames will exit so as to potentially

reduce frame traffic congestion along potential routes that include said set of possible exit ports comprises employing said multiple weights to select an exit port that is as good as or better than alternative exit ports in terms of achieving an objective reflected by said weights.

13. The method of claim 12, wherein as good as or better in terms of achieving an objective function comprises obtaining a higher value objective function.

14. The method of claim 12, wherein as good as or better in terms of achieving an objective function comprises obtaining a lower value objective function.

15. The method of claim 12, wherein said multiple weights at least in part reflect consumed bandwidth for particular routes.

16. The method of claim 8, wherein at least some of said links are to one or more other switches in a switch fabric.

17. The method of claim 1, wherein at least one of said set of possible exit ports is selected based at least in part on a source tag and/or a destination tag added to said frame after said frame enters said switch.

18. The method of claim 17, wherein said source tag and/or said destination tag is stripped off said frame before said frame exits said switch.

19. The method of claim 1, wherein at least one of said set of possible exist ports is selected based at least in part on a source tag and/or a destination tag added to each of said frames after said frames enter said switch.

20. The method of claim 19, wherein said source tag and/or said destination tag is stripped off each of said frames before each of said frames exits said switch.

21. A switch fabric comprising:

at least a first switch and a second switch;

said first and said second switch being communicatively coupled;

said first switch including a processor and a memory, and having the capability to balance a flow of frames exiting said first switch;

said first switch being adapted to select an exit port of said switch from a set of possible exit ports through which a frame from said flow of frames will exit so as to potentially reduce frame traffic congestion along potential routes that include said set of possible exit ports, said set of possible exit ports including at least some of the exit ports of at least two trunk groups.

22. The switch fabric of claim 21, wherein at least one of said trunk groups comprises four exit ports.

23. The switch fabric of claim 21, wherein at least one of said trunk groups comprises eight exit ports.

24. The switch fabric of claim 21, wherein said first switch is adapted to select said exit port pseudo-randomly.

25. The switch fabric of claim 24, wherein first said switch is adapted to select said exit port pseudo-randomly by applying a hash function.

26. The switch fabric of claim 25, wherein said first switch is adapted to apply said hash function to a set of parameters associated with the frames that will exit said first switch.

27. The switch fabric of claim 21, wherein said first switch is adapted to respectively assign a weight to at least some respective ones of said exit ports;

wherein said first switch is adapted to employ the weights to select an exit port for a frame of said flow to exit so as to balance said flow of frames exiting said first switch so that said selected exit port is as good as or better than alternative exit ports in terms of achieving an objective reflected by said weights.

28. The switch fabric of claim 27, wherein as good as or better in terms of achieving an objective function comprises obtaining a higher value objective function.

29. The switch fabric of claim 27, wherein as good as or better in terms of achieving an objective function comprises obtaining a lower value objective function.

30. The switch fabric of claim 27, wherein said weights are adapted to at least in part reflect consumed bandwidth for particular routes.

31. The switch fabric of claim 27, wherein at least some exit ports have multiple weights to reflect routes from the particular exit port to multiple respective destination ports;

wherein said first switch is adapted to employ said multiple weights to select an exit port for a frame of said flow to exit so as to balance said flow of frames exiting said first switch to select an exit port that is as good as or better than alternative exit ports in terms of achieving an objective reflected by said weights.

32. The switch fabric of claim 31, wherein said multiple weights are adapted to at least in part reflect consumed bandwidth for particular routes.

33. The switch fabric of claim 31, wherein as good as or better in terms of achieving an objective function comprises obtaining a higher value objective function.

34. The switch fabric of claim 31, wherein as good as or better in terms of achieving an objective function comprises obtaining a lower value objective function.

35. The switch fabric of claim 27, wherein at least some of said links are to one or more other switches in said switch fabric other than said second switch.

36. The switch fabric of claim 21, wherein said first switch is adapted to select at least one of said set of possible exit ports based at least in part on a source tag and/or a destination tag added to said frame after said frame enters said switch.

37. The switch fabric of claim 36, wherein said first switch is adapted to strip said source tag and/or said destination tag off said frame before said frame exits said switch.

38. The switch fabric of claim 21, wherein said first switch is adapted to select at least one of said set of possible exit ports based at least in part on a source tag and/or a destination tag added to each of said frames after said frames enter said switch.

39. The switch fabric of claim 38, wherein said first switch is adapted to strip said source tag and/or said destination tag off each of said frames before each of said frames exits said switch.

40. An apparatus comprising:

a switch, said switch including a processor and a memory;

said switch further having the capability to balance a flow of frames exiting said switch;

said switch being adapted to select an exit port of said switch from a set of possible exit ports through which a frame from said flow of frames will exit so as to potentially reduce frame traffic congestion along potential routes that include said set of possible exit ports, said set of possible exit ports including at least some of the exit ports of at least two trunk groups.

41. The apparatus of claim 40, wherein at least one of said trunk groups comprises four exit ports.

42. The apparatus of claim 40, wherein at least one of said trunk groups comprises eight exit ports.

43. The apparatus of claim 40, wherein said switch is adapted to select said exit port pseudo-randomly.

44. The apparatus of claim 43, wherein said switch is adapted to select said exit port pseudo-randomly by applying a hash function.

45. The apparatus of claim 44, wherein said switch is adapted to apply said hash function to a set of parameters associated with the frames that will exit said switch.

46. The apparatus of claim 40, wherein said switch is adapted to respectively assign a weight to at least some respective ones of said exit ports;

wherein said switch is adapted to employ the weights to select an exit port for a frame of said flow to exit so as to balance said flow of frames exiting said switch so that said selected exit port is as good as or better than alternative exit ports in terms of achieving an objective reflected by said weights.

47. The apparatus of claim 46, wherein as good as or better in terms of achieving an objective function comprises obtaining a lower value objective function.

48. The apparatus of claim 46, wherein as good as or better in terms of achieving an objective function comprises obtaining a higher value objective function.

49. The apparatus of claim 46, wherein said weights at least in part reflect consumed bandwidth for particular routes.

50. The apparatus of claim 46, wherein at least some exit ports have multiple weights to reflect routes from the particular exit port to multiple respective destination ports;

wherein said switch is adapted to employ said multiple weights to select an

exit port for a frame of said flow to exit so as to balance said flow of frames exiting said switch to select an exit port that is as good as or better than alternative exit ports in terms of achieving an objective reflected by said weights.

51. The apparatus of claim 50, wherein as good as or better in terms of achieving an objective function comprises obtaining a lower value objective function.

52. The apparatus of claim 50, wherein as good as or better in terms of achieving an objective function comprises obtaining a higher value objective function.

53. The apparatus of claim 50, wherein said multiple weights at least in part reflect consumed bandwidth for particular routes.

54. The apparatus of claim 46, wherein at least some of said links are to one or more other switches in a switch fabric.

55. The apparatus of claim 41, wherein said switch is adapted to select at least one of said set of possible exit ports based at least in part on a source tag and/or a destination tag added to said frame after said frame enters said switch.

56. The apparatus of claim 55, wherein said switch is adapted to strip said source tag and/or said destination tag off said frame before said frame exits said switch.

57. The apparatus of claim 40, wherein said switch is adapted to select at least one of said set of possible exit ports based at least in part on a source tag and/or a

destination tag added to each of said frames after said frames enter said switch.

58. The apparatus of claim 57, wherein said switch is adapted to strip said source tag and/or said destination tag off each of said frames before each of said frames exits said switch.

59. A network comprising:

a host;

a physical storage unit;

a first switch and a second switch communicatively coupled to form a switch fabric;

said first switch and said second switch further communicatively coupled to said host and said physical storage unit;

at least said first switch including a processor and memory, and having the capability to balance a flow of frames exiting said switch;

said first switch being adapted to select an exit port of said switch from a set of possible exit ports through which a frame from said flow of frames will exit so as to potentially reduce frame traffic congestion along potential routes that include said set of possible exit ports, said set of possible exit ports including at least some of the exit ports of at least two trunk groups.

60. The network of claim 59, wherein at least one of said trunk groups comprises four exit ports.

61. The network of claim 59, wherein at least one of said trunk groups comprises eight exit ports.

62. The network of claim 59, wherein said first switch is adapted to select said exit port pseudo-randomly.

63. The network of claim 62, wherein said first switch is adapted to select said exit port pseudo-randomly by applying a hash function.

64. The network of claim 63, wherein said first switch is adapted to apply said hash function to a set of parameters associated with the frames that will exit said first switch.

65. The network of claim 59, wherein said first switch is adapted to respectively assign a weight to at least some respective ones of said exit ports;

wherein said first switch is adapted to employ the weights to select an exit port for a frame of said flow to exit so as to balance said flow of frames exiting said first switch so that said selected exit port is as good as or better than alternative exit ports in terms of achieving an objective reflected by said weights.

66. The network of claim 65, wherein as good as or better in terms of achieving an objective function comprises obtaining a lower value objective function.

67. The network of claim 65, wherein as good as or better in terms of achieving an objective function comprises obtaining a higher value objective function.

68. The network of claim 65, wherein said weights at least in part reflect consumed bandwidth for particular routes.

69. The network of claim 65, wherein at least some exit ports have multiple weights to reflect routes from the particular exit port to multiple respective destination ports;

wherein said first switch is adapted to employ said multiple weights to select an exit port for a frame of said flow to exit so as to balance said flow of frames exiting said first switch to select an exit port that is as good as or better than alternative exit ports in terms of achieving an objective reflected by said weights.

70. The network of claim 69, wherein as good as or better in terms of achieving an objective function comprises obtaining a lower value objective function.

71. The network of claim 69, wherein as good as or better in terms of achieving an objective function comprises obtaining a higher value objective function.

72. The network of claim 69, wherein said multiple weights at least in part reflect consumed bandwidth for particular routes.

73. The network of claim 65, wherein at least some of said links are to one or more other switches in said switch fabric other than said second switch.

74. The network of claim 59, wherein said first switch is adapted to select at least one of said set of possible exit ports based at least in part on a source tag and/or a destination tag added to said frame after said frame enters said first switch.

75. The network of claim 74, wherein said first switch is adapted to strip said source tag and/or said destination tag off said frame before said frame exits said first switch.

76. The network of claim 59, wherein said first switch is adapted to select at least one of said set of possible exist ports based at least in part on a source tag and/or a destination tag added to each of said frames after said frames enter said first switch.

77. The network of claim 76, wherein said first switch is adapted to strip said source tag and/or said destination tag off each of said frames before each of said frames exits said first switch.

78. An article comprising: a storage medium having stored thereon instructions that, when executed, result in performance of a method of balancing a flow of frames exiting a switch that includes the following:

applying a process to select an exit port of said switch from a set of possible exit ports through which a frame from said flow of frames will exit so as to potentially reduce frame traffic congestion along potential routes that include said set of possible exit ports, said set of possible exit ports including at least some of the exit ports of at least two trunk groups.

79. The article of claim 78, wherein at least one of said trunk groups comprises four exit ports.

80. The article of claim 78, wherein at least one of said trunk groups comprises eight exit ports.

81. The article of claim 78, wherein said instructions, when executed, further

result in: said process comprising a pseudo-random process.

82. The article of claim 81, wherein said instructions, when executed, further result in: said applying said pseudo-random process comprises applying a hash function.

83. The article of claim 82, wherein said instructions, when executed, further result in: said hash function being applied to a set of parameters associated with the frames exiting said switch.

84. The article of claim 78, wherein said instructions, when executed, further result in: a weight being respectively assigned to at least some respective ones of said exit ports; and further result in: applying a process to select an exit port for a frame of said flow to exit so as to balance said flow of frames exiting said switch comprising employing the weights to select an exit port that is as good as or better than alternative exit ports in terms of achieving an objective reflected by said weights.

85. The article of claim 84, wherein said instructions, when executed, further result in: said weights at least in part reflecting consumed bandwidth for particular routes.

86. The article of claim 84, wherein said instructions, when executed, further result in: at least some exit ports having multiple weights reflecting routes from the particular exit port to multiple respective destination ports; and further result in: applying a process to select an exit port for a frame of said flow to exit so as to

balance said flow of frames exiting said switch comprising employing said multiple weights to select an exit port that is as good as or better than alternative exit ports in terms of achieving an objective reflected by said weights.

87. The article of claim 86, wherein said instructions, when executed, further resulting in: multiple weights at least in part reflecting consumed bandwidth for particular routes.

88. The article of claim 78, wherein said instructions, where further executed, result in: at least one of said set of possible exit ports being selected based at least in part on a source tag and/or a destination tag added to said frame after said frame enters said switch.

89. The article of claim 88, wherein said instructions, when further executed, result in: said source tag and/or said destination tag being stripped off said frame before said frame exits said switch.

90. The article of claim 78, wherein said instructions, when further executed, result in: at least one of said set of possible exist ports being selected based at least in part on a source tag and/or a destination tag added to each of said frames after said frames enter said switch.

91. The article of claim 90, wherein said instructions, when further executed, result in: said source tag and/or said destination tag being stripped off each of said frames before each of said frames exits said switch.

92. An article comprising: a storage medium having stored thereon instructions

that, when executed, result of a method of initializing a switch to route a flow of frames comprising:

initializing said switch to apply a process to select an exit port of said switch from a set of possible exit ports through which a frame from said flow of frames will exit so as to potentially reduce frame traffic congestion along potential routes that include said set of possible exit ports, said set of possible exit ports including at least some of the exit ports of at least two trunk groups.

93. The article of claim 92, wherein said instructions, when executed, further result in: said switching being initialized to apply a pseudo-random process.

94. The article of claim 93, wherein said instructions, when executed, further result in: said switch being initialized to apply a hash function being to a set of parameters associated with frames exiting said switch.

95. The article of claim 92, wherein said instructions, when executed, further result in: said switch being initialized to assign a weight to at least some respective ones of said exit ports; and said switch being initialized to apply a process to select an exit port for a frame of said flow to exit so as to balance said flow of frames exiting said switch comprising employing the weights to select an exit port that is as good as or better than alternative exit ports in terms of achieving an objective reflected by said weights.